arXiv:1510.03895v1 [cs.DS] 13 Oct 2015

A FASTER SUBQUADRATIC ALGORITHM FOR FINDING OUTLIER CORRELATIONS

MATTI KARPPA, PETTERI KASKI, AND JUKKA KOHONEN

Abstract. We study the problem of detecting *outlier pairs* of strongly correlated variables among a collection of n variables with otherwise weak pairwise correlations. After normalization, this task amounts to the geometric task where we are given as input a set of n vectors with unit Euclidean norm and dimension d, and we are asked to find all the outlier pairs of vectors whose inner product is at least ρ in absolute value, subject to the promise that all but at most q pairs of vectors have inner product at most τ in absolute value for some constants $0 < \tau < \rho < 1$.

Improving on an algorithm of G. Valiant [FOCS 2012; J. ACM 2015], we present a randomized algorithm that for Boolean inputs ($\{-1, 1\}$-valued data normalized to unit Euclidean length) runs in time

$$\tilde{O}\big(n^{\max\{1-\gamma+M(\Delta\gamma,\gamma),\, M(1-\gamma,2\Delta\gamma)\}} + qdn^{2\gamma}\big)\,,$$

where $0 < \gamma < 1$ is a constant tradeoff parameter and $M(\mu, \nu)$ is the exponent to multiply an $\lfloor n^\mu \rfloor \times \lfloor n^\nu \rfloor$ matrix with an $\lfloor n^\nu \rfloor \times \lfloor n^\mu \rfloor$ matrix and $\Delta = 1/(1 - \log_\tau \rho)$. As corollaries we obtain randomized algorithms that run in time

$$\tilde{O}\big(n^{\frac{2\omega}{3-\log_\tau \rho}} + qdn^{\frac{2(1-\log_\tau \rho)}{3-\log_\tau \rho}}\big)$$

and in time

$$\tilde{O}\big(n^{\frac{4}{2+\alpha(1-\log_\tau \rho)}} + qdn^{\frac{2\alpha(1-\log_\tau \rho)}{2+\alpha(1-\log_\tau \rho)}}\big)\,,$$

where $2 \le \omega < 2.38$ is the exponent for square matrix multiplication and $0.3 < \alpha \le 1$ is the exponent for rectangular matrix multiplication. We present further corollaries for the light bulb problem and for learning sparse Boolean functions. (The notation $\tilde{O}(\cdot)$ hides polylogarithmic factors in n and d whose degree may depend on ρ and τ.)

1. Introduction

1.1. Scalable correlation analysis.

The *Pearson product-moment correlation coefficient* [20, 46, 51], or briefly, *correlation*, is one of the fundamental statistical quantities used to measure the strength of interaction between two sequences of d numerical observations. Assuming we have n such sequences, there are $\Theta(n^2)$ pairs of sequences, which immediately puts forth the algorithmic question of *scalability* of correlation-based analyses as our number of observables n increases.

If we seek to explicitly compute all $\Theta(n^2)$ pairwise correlations, the size of our desired result forces us to quadratic $\Theta(n^2)$ scaling as a function of n. Thus, assuming we seek *subquadratic*[1] scaling in n, we must refine our objective towards aggregate analyses that only implicitly consider all the pairwise correlations. One natural goal in this setting is to identify *interesting* pairs

[1]For brevity, we will use the expression "subquadratic in n" to mean $O(n^{2-\epsilon})$ for some constant $\epsilon > 0$ independent of n.

of observables, subject to the assumption that there are few such pairs. In the context of correlation, one possible signal for an interesting interaction between a pair of observables is an *abnormally large*, or *outlier*, (absolute) correlation, as measured against a *background* parameter that bounds from above the absolute values of correlations between most pairs of observables.

Assuming our n observables are each normalized to zero mean and unit standard deviation, the computational task of finding abnormally large correlations reduces to the following geometric setup:

Problem 1 (Outlier Correlations)**.** Given as input a set of n vectors with unit Euclidean norm and dimension d, find all the *outlier pairs* of vectors whose inner product is at least ρ in absolute value, subject to the *promise* that *all but at most q pairs* of vectors have inner product at most τ in absolute value, for some $0 < \tau < \rho < 1$.

(The promise implies in particular that there are at most q outlier pairs. Our interest is on inputs where q is subquadratic in n.)

Remark 2. A useful variant of the problem is that the input vectors are each assigned one of two possible colors, and we want to find all the outlier pairs with distinct colors. Let us call this variant Bichromatic Outlier Correlations.

The question of scalability as n increases is now parameterized by the dimension d of the vectors and the thresholds ρ, τ. Furthermore, we can impose restrictions on the structure of the vectors themselves to study restricted variants of the problem. From a discrete computational standpoint, a natural restriction is to assume the input vectors take values in $\{-1, 1\}$.[2] (Indeed, two vectors in $x, y \in \{-1, 1\}^d$ have inner product $\langle x, y\rangle = \sum_{i=1}^{d} x_i y_i = \lambda d$ if and only if the corresponding normalized vectors have inner product $\langle x/\sqrt{d}, y/\sqrt{d}\rangle = \lambda$ with $-1 \leq \lambda \leq 1$.)

For brevity, let us call $\{-1, 1\}$-valued inputs *Boolean* inputs.

Remark 3. For Boolean vectors $x, y \in \{-1, 1\}^d$, the inner product and the *Hamming distance* $D_H(x, y) = |\{i \in \{1, 2, \ldots, d\} : x_i \neq y_i\}|$ are affinely equivalent via the identity

$$\langle x, y\rangle + 2D_H(x, y) = d\,. \tag{1}$$

Thus, for Boolean inputs the Outlier Correlations problem is a variant of *near neighbor search* under the Hamming metric—the larger the inner product, the smaller the Hamming distance, and vice versa. (Cf. [26, 38, 65].)

1.2. **Earlier lower and upper bounds.** Already for Boolean inputs it is known that a quest for subquadratic scaling in n is faced with fine-grained complexity-theoretic barriers. Indeed, unless the Strong Exponential Time Hypothesis [24] is false, any algorithm that for a bichromatic input detects the existence of an outlier pair with $n^{2-\epsilon}$ scaling in n (and simultaneous $2^{o(d)}$ scaling in the dimension d) must have the property that ϵ depends on ρ and

[2] An arbitrary input on the unit Euclidean sphere can be quantized to a $\{-1, 1\}$-valued input if one is willing to accept randomization and an expected loss $\lambda \mapsto 1 - \frac{2}{\pi} \arccos \lambda$ in the parameters $\lambda \in \{\tau, \rho\}$ caused by rounding the arbitrary input into a Boolean input with random hyperplanes. See e.g. [53, Algorithm 4] or [11].

τ. More precisely, it follows by a simple local transformation (see §6) from a result of Williams [59] that unless SETH is false, ϵ cannot be bounded from below by a positive constant when $|\rho - \tau| \to 0$.

On the positive side, algorithms for the $(1+\epsilon)$-*approximate nearest neighbor* problem can be used to identify outlier pairs.[3] *Locality-sensitive hashing* (LSH), introduced by Indyk and Motwani [27], is a well-known framework for finding $(1+\epsilon)$-approximate nearest neighbors. In their seminal paper, Indyk and Motwani describe an algorithm that, for $\epsilon > 0$ and the Hamming metric, uses $\tilde{O}(n^{1+1/(1+\epsilon)} + dn)$ space for supporting approximate nearest-neighbor queries to a data set of n vectors of dimension d, and requires $\tilde{O}(dn^{1/(1+\epsilon)})$ query time.[4] Lower bounds are known for the hash function families [40, 42] which guarantee that algorithms based on data-independent LSH, such as Andoni and Indyk [5] with query time $dn^{1/(1+\epsilon)^2+o(1)}$ for the Euclidean metric, are essentially optimal. Andoni, Indyk, Nguyen and Razenshteyn [7] present a framework that circumvents these lower bounds by introducing *data dependence* into the hash families, giving the improved query exponent $7/(8(1+\epsilon)^2) + O(1/(1+\epsilon)^3) + o(1)$ for the query time in the Euclidean metric. Recently, Andoni and Razenshteyn [8] optimize this exponent to $1/(2(1+\epsilon)^2-1)+o(1)$ and Andoni *et al.* present LSH optimized for *angular distance* [6].

Turning from approximate to *exact* solutions, Alman and Williams [4] show that the *batch* Hamming nearest neighbor problem with n simultaneous queries to a database of n vectors can be solved in randomized $n^{2-1/O(c(n)\log c(n))}$ time if $d = c(n)\log n$ for $c(n) : \mathbb{N} \to \mathbb{N}$. This scaling is subquadratic in n if the dimension d is bounded by $d = O(\log n)$.

1.3. **The curse of weak outliers and Valiant's breakthrough.** For Boolean inputs with unrestricted dimension, the scaling obtained via approximate nearest neighbors is subquaradtic in n for all constant $\rho > \tau$. However, such scaling suffers from what could be called the *curse of weak outliers*. Namely, if ρ is small (that is, the outlier correlations themselves are weak), the scaling in n is essentially quadratic $n^{2-c\rho}$ for a positive constant c, *even if the background* τ *decays to zero*. Ideally, there should be a way to avoid such direct dependence on the value of ρ, as long as ρ and τ do not converge.

In a breakthrough result, G. Valiant [53] showed that subquadratic scaling *is* possible, *independently of* ρ, as long as $\log_\tau \rho = (\log \rho)/(\log \tau)$ remains bounded from above by a small positive constant.

---

[3] By (1), two Boolean vectors with inner product λd have Hamming distance $(1-\lambda)d/2$. Thus, assuming there are no pairs whose inner product is strictly between τ and ρ, the vectors that belong to outlier pairs can be discovered by selecting a small enough ϵ so that an $(1+\epsilon)$-approximate nearest neighbor of a vector has to be a completion of such a vector to an outlier pair if such a completion exists. For example, by (1) we can take $1+\epsilon < (1-\tau)/(1-\rho)$ for the Hamming metric and $(1+\epsilon)^2 < (1-\tau)/(1-\rho)$ for the Euclidean metric.

[4] The notation $\tilde{O}(\cdot)$ suppresses factors polylogarithmic in n and d whose degree may depend on ρ and τ.

Let $2 \leq \omega < 2.3728639$ be the exponent[5] of square matrix multiplication [32]. We say that a parameterized event E_n happens *with high probability (w.h.p.) in the parameter* n if $\Pr(E_n) \geq 1 - f(n)$ for a function $f(n)$ with $\lim_{n\to\infty} n^k f(n) = 0$ for all positive k.

Theorem 4 (Valiant [53])**.** *For all constants* $0 < \tau < \rho < 1$*, the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time* $\tilde{O}(n^{\frac{5-\omega}{4-\omega}+\omega \log_\tau \rho})$ *for approximate detection*[6] *and subsequent* $\tilde{O}(qdn^{1/(4-\omega)})$ *time for exact listing of all the outliers, w.h.p.*

Remark 5. A tacit property of Valiant's algorithm is that the running time for approximate detection is, up to polylogarithmic factors, *independent of the dimension* d of the input. In particular, since the input is Boolean, the (normalized) input vectors are *decoherent*[7] and hence it suffices to randomly sample the input to access the correlations.

Valiant's result opens up a quest to understand the extent of subquadratic scaling available for Outlier Correlations:

Assuming that the outliers are well-separated from the background correlations, that is, that $\log_\tau \rho$ *is small, how close to linear scaling in* n *can we get?*

Our intent in this paper is to present further progress in such a quest.

1.4. **Valiant's algorithm in more detail.** Since our intent is to improve on Valiant's algorithm, it will be convenient to review the key ideas in its design. Suppose our input data for Bichromatic[8] Outlier Correlations is given to us as two $d \times n$ Boolean matrices, $A \in \{-1, 1\}^{d\times n}$ and $B \in \{-1, 1\}^{d\times n}$. Our task is to find the outlier inner products (of absolute value at least ρd) between columns of A and B. Valiant's algorithm proceeds in the following four phases; the first and second phases together constitute

[5]We adopt the convention that all running time bounds that contain an exponent $f(\omega)$ of n that depends on ω (or any other limiting exponent for matrix multiplication) are tacitly stated as $f(\omega) + \epsilon$ for an arbitrarily small constant $\epsilon > 0$.

[6]We may stop the algorithm after the approximate detection phase if we only want to decide whether the input contains at least one pair with absolute inner product at least ρd, with the following *approximation* guarantees: (i) if all the inner products in the input have absolute value at most τd, the algorithm outputs false w.h.p.; and (ii) if the input contains at least one inner product with absolute value at least ρd, the algorithm outputs true w.h.p.

[7]By *decoherence* we mean here the property that if we view the squares of the entries of a (normalized) input vector as a probability distribution over the d dimensions, the probability mass is roughly uniformly distributed across the dimensions. With (normalized) Boolean input, the distribution is *exactly* uniform and thus eases concentration analyses.

[8]Up to polylogarithmic factors in n, we can without loss of generality work with the bichromatic variant by a routine reduction. Indeed, suppose the data matrix for Outlier Correlations is $D \in \{-1, 1\}^{d\times n}$ and we are interested in discovering the pairs of distinct columns in D with (absolute) inner product at least ρd. Construct $\lceil \log_2 n \rceil$ pairs of matrices A_k, B_k for $k = 0, 1, \ldots, \lceil \log_2 n \rceil - 1$, where the matrix A_k (respectively, B_k) contains the columns of D whose index, viewed as a $\lceil \log_2 n \rceil$-bit string, has a 0 (respectively, 1) in bit position k. Take the union of the solutions of the instances A_k, B_k to obtain the solution for D. All distinct columns will be discovered because each pair of distinct columns differs in at least one bit position k.

a *compression* phase which enables the use of fast matrix multiplication in the third *detection* phase, which is followed by a *listing* phase.

1. Expansion by uniform random sampling of a tensor power. It is possible to *amplify* the inner products between columns of A and B by *individually* tensoring the matrices A and B, at the cost of increasing the dimension of the data from d to d^p for a positive integer p. Indeed, for any two d-dimensional real vectors it holds that $\langle x, y\rangle^p = \langle x^{\otimes p}, y^{\otimes p}\rangle$, where $\otimes p$ indicates p-fold Kronecker product (tensoring) of the vector with itself. Accordingly, taking p-fold Kronecker products *in the vertical dimension only* (which we indicate with the notation $\downarrow\otimes p$) we obtain the $d^p \times n$ matrices $A^{\downarrow\otimes p}$ and $B^{\downarrow\otimes p}$. A key observation is that a uniform random *sample* of size s of the d^p dimensions suffices *because the input is decoherent* (Boolean) and hence the sum of the sample is strongly concentrated, e.g., via the Hoeffding bounds [23]. Accordingly, we assume that the matrices $A^{\downarrow\otimes p}$ and $B^{\downarrow\otimes p}$ in fact have dimensions $s \times n$.

2. Signed aggregation. Since comparatively few of the n^2 inner products between the columns of $A^{\downarrow\otimes p}$ and $B^{\downarrow\otimes p}$ are outliers, after amplification it is possible to *aggregate* the n columns in $A^{\downarrow\otimes p}$ (respectively, $B^{\downarrow\otimes p}$) by randomly partitioning the columns into n/t *blocks* of t columns, and taking, with an independent uniform random sign for each column, the sum of the columns in each block. This produces an $s \times (n/t)$ matrix $\tilde{A}$ (respectively, $\tilde{B}$).

3. Detection. Because the outer dimension has now decreased from n to n/t, with careful selection of the parameters s, t, p we can now afford to multiply the compressed matrices $\tilde{A}, \tilde{B}$ in subquadratic time to obtain the $(n/t) \times (n/t)$ product $\tilde{A}^\top \tilde{B}$. Because of careful amplification by sampling and signed aggregation, the pairs of blocks containing at least one outlier pair are with moderate probability signalled by entries in $\tilde{A}^\top \tilde{B}$ that have absolute value above a threshold value.[9]

4. Listing. Finally, the outlier pairs can be computed with brute force by computing the t^2 pairwise inner products of columns of A and B within each signalled pair of blocks.

Theorem 4 follows by careful selection of the parameters s, t, p and the signalling threshold.

1.5. **Our contribution.** Our contribution in this paper amounts to the observation that the compression phase (expansion followed by signed aggregation) in Valiant's algorithm can be replaced by a faster compression subroutine. In essence, we rely on *fast matrix multiplication* as the algorithmic device to *simultaneously* expand and aggregate; this requires that we replace the uniform random sampling of dimensions in the expansion phase of Valiant's algorithm with *Cartesian product sampling* to enable us to entangle expansion and aggregation. Despite the considerable decrease in entropy compared with a uniform random sample, we show that Cartesian product sampling *on a tensor power* remains roughly as sharply concentrated as a uniform random sample, enabling us to use roughly the same sample size s as Valiant's algorithm for comparable compression, thus resulting in faster

[9] The algorithm is iterated $\Theta(\log n)$ times to guarantee that every outlier pair is signalled in at least one pair of blocks during at least one iteration with high probability.

execution because of speedup given by fast matrix multiplication. The faster compression subroutine then enables a faster tradeoff that balances between compression and detection, as controlled by a tradeoff parameter $0 < \gamma < 1$.

Let us now state our main result for Outlier Correlations. For constants $\mu, \nu > 0$, let $M(\mu, \nu)$ be the infimum of the values $\sigma > 0$ such that there exists an algorithm that multiplies an $\lfloor n^\mu \rfloor \times \lfloor n^\nu \rfloor$ integer matrix with an $\lfloor n^\nu \rfloor \times \lfloor n^\mu \rfloor$ integer matrix in $O(n^\sigma)$ arithmetic operations. For example, $\omega = M(1, 1)$.

Theorem 6 (Main, basic tradeoff)**.** *For all constants $0 < \gamma < 1$ and $\Delta \geq 1$, the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time*

$$\tilde{O}\big(n^{\max\{1-\gamma+M(\Delta\gamma,\gamma),\, M(1-\gamma,2\Delta\gamma)\}}\tau^{-4}\big) \tag{2}$$

for approximate detection and subsequent

$$\tilde{O}\big(qdn^{2\gamma}\big)$$

time for exact listing of all the outliers, w.h.p. The running time bounds hold uniformly for all $n^{-\Theta(1)} \leq \tau < \rho < 1$ with $\log_\tau \rho \leq 1 - \Delta^{-1}$.

For specific choices of γ and Δ we obtain the following corollaries.

Corollary 7. *For all constants $0 < \tau < \rho < 1$, the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time*

$$\tilde{O}\big(n^{\frac{2\omega}{3-\log_\tau \rho}}\big)$$

for approximate detection and subsequent

$$\tilde{O}\big(qdn^{\frac{2(1-\log_\tau \rho)}{3-\log_\tau \rho}}\big)$$

time for exact listing of all the outliers, w.h.p.

Let $0.30298 < \alpha \leq 1$ be the exponent for rectangular matrix multiplication [31]. That is, α is the supremum of all values $\sigma \leq 1$ with $M(1, \sigma) = 2$.

Corollary 8. *For all constants $0 < \tau < \rho < 1$, the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time*

$$\tilde{O}\big(n^{\frac{4}{2+\alpha(1-\log_\tau \rho)}}\big)$$

for approximate detection and subsequent

$$\tilde{O}\big(qdn^{\frac{2\alpha(1-\log_\tau \rho)}{2+\alpha(1-\log_\tau \rho)}}\big)$$

time for exact listing of all the outliers, w.h.p.

Figure 1 displays the subquadratic scaling in n for approximate detection obtained from Corollaries 7 and 8, contrasted with Theorem 4 of Valiant [53].

Remark 9. Contrasted with Theorem 4, Corollary 8 implies that asymptotically it is possible to list extremely weak outliers with, say, $\rho = 2^{-100}$ and $\tau = 2^{-101}$, in time $\tilde{O}(n^{1.998} + qdn^{0.003})$.

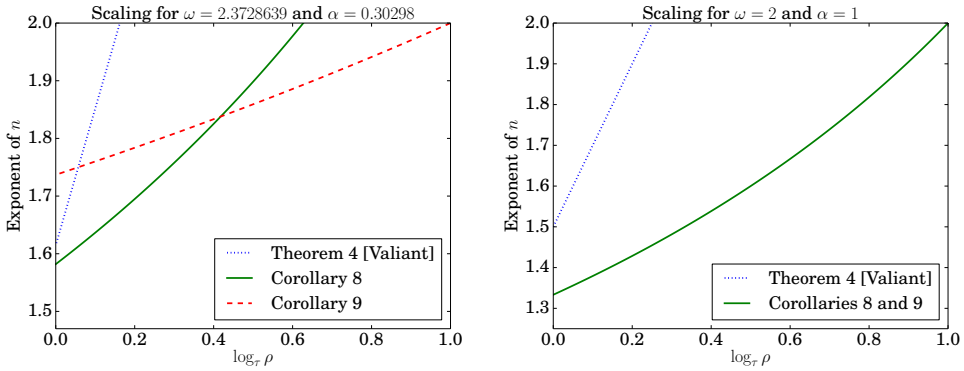

FIGURE 1. Illustration of subquadratic scaling in n for approximate detection as a function of the parameter $\log_\tau \rho$. The matrix multiplication exponents ω and α are fixed to values $\omega = 2.3728639$ and $\alpha = 0.30298$ (left), and to values $\omega = 2$ and $\alpha = 1$ (right).

Remark 10. In the setting of well-separated outliers with any constant ρ and $\tau \to 0$, the exponent for approximate detection in Corollary 7 approaches $2\omega/3$, improving Valiant's exponent $(5-\omega)/(4-\omega)$ across the range $2 \leq \omega < 2.38$.

Although our results show that exponents as low as $2\omega/3 = 4/3$ may be feasible for finding outlier correlations (if we assume $\omega = 2$ and let $\tau \to 0$), the present framework unfortunately appears not to be powerful enough to lower the exponent below $4/3$. It remains open whether the exponent can be lowered all the way to 1 (to linear scaling in n), and whether techniques other than fast matrix multiplication can be used to attain subquadratic scaling without the curse of weak outliers.

2. Related work and applications

The Outlier Correlations problem admits many generalizations and applications. Beyond the direct application of finding correlated observables in a statistical context, it is a standard practice in, among others, machine learning and data analysis to model the objects of interest as d-dimensional vectors, whereby the *inner product* (normalized or unnormalized with the norms of the two vectors) is a very popular measure for *similarity* among the underlying objects (cf. [2, 9, 11, 12, 14, 15, 16, 29, 33, 34, 41, 47, 48, 49, 50, 52, 55, 56, 57, 62, 63, 64]). Here we will be content on discussing a narrow set of applications of our results and related work. The proofs of all corollaries appear in §5.

2.1. The light bulb problem.

Theorem 6 gives as an almost immediate corollary a faster subquadratic algorithm (cf. Valiant [53, Corollary 2.2]) for solving the *light bulb problem* [54]. Let us first state the problem and then give our improvement.

Problem 11 (Light Bulb)**.** Suppose we are given as input a set of n vectors in $\{-1,1\}^d$ consisting of (i) a *planted pair* with inner product at

least ρd in absolute value for $0 < \rho < 1$, and (ii) $n-2$ independent uniform random vectors in $\{-1,1\}^d$. Our task is to find the planted pair among the n vectors.[10]

Corollary 12. *For any constant $0 < \epsilon < \omega/3$, the* Light Bulb *problem admits a randomized algorithm that for any*

$$d \geq 5\rho^{-\frac{4\omega}{9\epsilon}-\frac{2}{3}} \log n$$

runs in time

$$\tilde{O}\big(n^{\frac{2\omega}{3}+\epsilon}\rho^{-\frac{8\omega}{9\epsilon}-\frac{4}{3}}\big) \tag{3}$$

and finds the planted pair of vectors, with probability $1-o(1)$. The running time bound holds uniformly for all $n^{-\Theta(1)} < \rho < 1$.

Prior to Valiant's algorithm, the first subquadratic algorithm for the light bulb problem was the randomized $\tilde{O}(n^{1+(\log\frac{1+\rho}{2})/(\log\frac{1+\sigma}{2})})$ time algorithm of Paturi, Rajasekaran, and Reif [45], where σ measures the absolute value of the inner product between the pair of vectors with the second largest inner product. Besides the algorithms based on approximate nearest neighbors and locality-sensitive hashing (see §1.2) that can also be used to solve light bulb problem, the *bucketing codes* approach of Dubiner [18] yields a randomized $\tilde{O}(n^{2/(\rho+1)})$ time algorithm for the light bulb problem, which was the fastest algorithm prior to Valiant's algorithm. Recently, May and Ozerov [36] present a recursive approach for solving the light bulb problem; with weak outliers also their algorithm converges to quadratic running time.

2.2. **Learning Boolean functions.** The light bulb problem generalizes to the task of *learning a parity function in the presence of noise*.

Problem 13 (Parity with Noise)**.** Let the *support* $S \subseteq \{1,2,\ldots,n\}$ of the parity function have size $k = |S|$ and let the *noise rate* be $0 < \eta < 1$. Our task is to find S, given access to independent *examples* of the form (x,y), where (i) the *input* $x \in \{-1,1\}^n$ is chosen uniformly at random, and (ii) the *label* $y = z \cdot \prod_{j\in S} x_j$ is defined by independently choosing $z \in \{-1,1\}$ with $\Pr(z=-1) = \eta$.

Blum, Kalai, and Wasserman [10] and Lyubashevsky [35] study the general case of unrestricted k. Studies of the *sparse* case with $k = O(1)$ include works by Grigorescu, Reyzin, and Vempala [22] and Valiant [53]. In the sparse case, we can use a split-and-list transformation presented by Valiant [53, p. 32] together with the algorithm underlying Corollary 12 to essentially match an algorithm of Valiant that relies on Fourier-analytic techniques (cf. Valiant [53, Theorems 2.4, 5.2, and 5.5]). Compared with Valiant's algorithm, our algorithm has a slightly worse tolerance for noise but better sample complexity:

[10] It follows e.g. from the Hoeffding bound (6) and the union bound that there is a constant $c > 0$ such that for $d \geq c\rho^{-2}\log n$ with probability $1-o(1)$ the planted pair is the unique pair of vectors in the input with inner product at least ρd in absolute value.

Corollary 14. *For any constant $0 < \epsilon < \omega/3$, the* Parity with Noise *problem admits a randomized algorithm that uses*

$$d \geq (2k+3)|1-2\eta|^{-\frac{4\omega}{9\epsilon}-\frac{2}{3}} \log n \tag{4}$$

examples, for all sufficiently large k runs in time

$$\tilde{O}\big(n^{\frac{\omega+\epsilon}{3}k}|1-2\eta|^{-\frac{8\omega}{9\epsilon}-\frac{4}{3}}\big)\,, \tag{5}$$

and finds the support of the parity function, with probability[11] *at least $1/4 - o(1)$. The running time bound holds uniformly for all $n^{-\Theta(1)} < |1-2\eta| < 1$.*

Using algorithms of Feldman, Gopalan, Khot, and Ponnuswami [19], and Mossel, O'Donnell, and Servedio [39], from Corollary 14 one can obtain further corollaries in the context of learning sparse juntas and DNFs. We refer to Valiant [53] for a detailed exposition.

2.3. **Subquadratic algorithms in small dimension.** If the dimension d is very small, then subquadratic algorithms for Outlier Correlations are available through dedicated space-partitioning data structures such as *Voronoi diagrams*, enabling near neighbor query times $\tilde{O}(d^c)$ for a constant c, but with exponential scaling in size as a function of d [37, 58]. Similar query-vs-size tradeoffs are available also in higher dimensions; for example, Kushilevitz, Ostrovsky, and Rabani [30] presented a data structure with $\tilde{O}(d)$ query time and size $(dn)^{O(1)}$ for approximate nearest neighbors. Recently, Alman and Williams [4] obtain subquadratic scaling for $d = O(\log n)$ (see §1.2).

2.4. **Scaling in the number q of above-background-correlated pairs.** Let us briefly study the case when q is relatively large, such as the *batch-query*[12] case $q = n$, or the case $q = O(n^{2-\delta})$ for a small constant $\delta > 0$. In this situation we observe that in Theorem 6 the term $\tilde{O}(qdn^{2\gamma})$ can dominate the running time over (2). Indeed, assuming qd is subquadratic in n, we can always obtain an overall running time that is subquadratic in n by forcing a small enough $\gamma > 0$ in Theorem 6. However, such a forced γ may cause a suboptimal value of (2) over what would be available if q was smaller. This bottleneck in scaling for large q can be somewhat alleviated by pursuing a two-level recursive strategy, where we run the approximate detection algorithm recursively before proceeding to listing:

Theorem 15 (Main, two-level recursive tradeoff)**.** *For all constants $0 < \gamma, \kappa < 1$ and $\Delta \geq 1$, the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time*

$$\tilde{O}\big(n^{\max\{1-\gamma+M(\Delta\gamma,\gamma),\, M(1-\gamma,2\Delta\gamma)\}}\tau^{-4}\big)$$

for approximate detection and subsequent

$$\tilde{O}\big(qn^{\gamma\max\{1-\kappa+M(\Delta\kappa,\kappa),\, M(1-\kappa,2\Delta\kappa)\}}\tau^{-4} + qdn^{2\gamma\kappa}\big)$$

[11] This success probability can be boosted by independent repetition by multiplying (4) and (5) with the number of repetitions; cf. §5.4.

[12] Suppose we have n observables in a database and n observables that constitute queries to the database. Suppose furthermore that for each query there are $O(1)$ outlier-correlated observables in the database, and that our task is to find these outlier-correlated observables for each query.

time for exact listing of all the outliers, w.h.p. The running time bounds hold uniformly for all $n^{-\Theta(1)} \leq \tau < \rho < 1$ *with* $\log_\tau \rho \leq 1 - \Delta^{-1}$.

For specific choices γ, κ, Δ we obtain the analogs of Corollary 7 and 8.

Corollary 16. *For all constants* $0 < \tau < \rho < 1$, *the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time*

$$\tilde{O}\big(n^{\frac{2\omega}{3-\log_\tau \rho}}\big)$$

for approximate detection and subsequent

$$\tilde{O}\big(qn^{\frac{2\omega(1-\log_\tau \rho)}{(3-\log_\tau \rho)^2}} + qdn^{\frac{2(1-\log_\tau \rho)^2}{(3-\log_\tau \rho)^2}}\big)$$

time for exact listing of all the outliers, w.h.p.

Corollary 17. *For all constants* $0 < \tau < \rho < 1$, *the* Outlier Correlations *problem for Boolean inputs admits a randomized algorithm that runs in time*

$$\tilde{O}\big(n^{\frac{4}{2+\alpha(1-\log_\tau \rho)}}\big)$$

for approximate detection and subsequent

$$\tilde{O}\big(qn^{\frac{4\alpha(1-\log_\tau \rho)}{(2+\alpha(1-\log_\tau \rho))^2}} + qdn^{\frac{2\alpha^2(1-\log_\tau \rho)^2}{(2+\alpha(1-\log_\tau \rho))^2}}\big)$$

time for exact listing of all the outliers, w.h.p.

Remark 18. Recall that Corollary 8 lists outliers with $\rho = 2^{-100}$ and $\tau = 2^{-101}$ in time $\tilde{O}(n^{1.998} + qdn^{0.003})$. Corollary 17 improves this to time $\tilde{O}(n^{1.998} + qn^{0.003} + qdn^{0.0000045})$.

Remark 19. In the setting of well-separated outliers with any constant ρ and $\tau \to 0$, Corollary 7 lists outliers in time $\tilde{O}(n^{2\omega/3+\epsilon_\tau} + qdn^{2/3+\epsilon_\tau})$. Corollary 16 improves this to time $\tilde{O}(n^{2\omega/3+\epsilon_\tau} + qn^{2\omega/9+\epsilon_\tau} + qdn^{2/9+\epsilon_\tau})$. Here $\epsilon_\tau \to 0$ as $\tau \to 0$.

2.5. **Related work in compressed matrix multiplication.** In addition to Valiant's algorithm, another recent breakthrough in the context of finding and approximating the outlier entries in a matrix product $A^\top B$ of two given $d \times n$ matrices A, B was made by Pagh [43]. Pagh's algorithm also computes a *compressed version* of the product $A^\top B$ by using 2-wise independent hash families that are algebraically compatible with a cyclic group (of order b) to aggregate the operands A, B with random signs into preimages of the hash function along the outer dimension (n) and with explicit (uncompressed) summation along the inner dimension (d). Because of the compatibility with the cyclic group, the compressed operands can be multiplied in essentially linear time via cyclic convolution (using the FFT for the cyclic group). The compressed product gives unbiased estimates of the product entries $(A^\top B)_{j_1 j_2}$ with 2-wise independence controlling the variance via the Frobenius norm at $\|A^\top B\|_F^2/b$. Despite the essentially linear-time performance given by the use of FFT in evaluating the compressed products (thus enabling a large value of b), the combination of Frobenius control on variance and the linear scaling along the inner dimension d appears not to yield improvements for the light bulb problem or for Outlier Correlations.

2.6. **Discussion and further work.** To our knowledge, there are no lower bounds that would preclude the use of other multilinear algorithms than fast matrix multiplication (such as Pagh's brilliant use of fast cyclic convolution [43]) for Outlier Correlations, yet fast matrix multiplication remains the only known tool to obtain truly subquadratic scaling for weak outliers. As Valiant highlights [53, §6], it would be of considerable interest to avoid fast matrix multiplication altogether to obtain *practical* subquadratic scaling for moderately-sized n. In fact, from this perspective our present results arguably proceed in the wrong direction by making *heavier* use of fast matrix multiplication to obtain *asymptotically* faster subquadratic scaling.

Very recently, Ahle, Pagh, Razenshteyn, and Silvestri [3] show that inner product similarity joins are hard to approximate in subquadratic time in n unless both SETH and the Orthogonal Vectors Conjecture (see §6) are false. Stated in terms of Outlier Correlations, they show that Boolean inputs with $\log_\tau \rho = 1 - o(1/\sqrt{\log n})$ do not admit subquadratic scaling in n unless both SETH and OVC are false. Narrowing the gap between subquadratic approximability and inapproximability remains a topic for further work.

Let us conclude with a few questions for further work. Is $\log_\tau \rho$ the natural parameter for subquadratic scaling in the context of Outlier Correlations? Is it possible to improve the limiting exponent $2\omega/3$ in Corollary 7 and in Corollary 12? Curiously, the algorithm in our Corollary 14 and Valiant's algorithm [53, Algorithm 11] for Parity with Noise rely on somewhat different techniques, yet both arrive at the constant $\omega/3$ for learning sparse parities—is this just a coincidence? Finally, is it possible to derandomize our algorithms or improve their space usage? We refer to Pagh [44] and Kapralov [28] for recent work and motivation for derandomization and resource tradeoffs in the context of similarity search.

3. Preliminaries

This section collects terminology, notation, and background results used in the subsequent technical development.

3.1. **The Hoeffding bound.** The Hoeffding bound establishes sharp concentration around the expectation of a sum of independent terms assuming we have control on the support of the terms.

Theorem 20 (Hoeffding [23, Theorem 2])**.** *Let $Z_1, Z_2, \ldots, Z_s$ be independent random variables with $\ell_i \leq Z_i \leq u_i$ for all $i = 1, 2, \ldots, s$ and let $Z = Z_1 + Z_2 + \ldots + Z_s$. Then, for all $d > 0$ it holds that*

$$\Pr\big(Z - \mathrm{E}\big[Z\big] \geq d\big) \leq \exp\left(-\frac{2d^2}{\sum_{i=1}^{s}(u_i - \ell_i)^2}\right). \tag{6}$$

3.2. **Inner products, restriction, amplification by powering.** Let us set up some notation regarding inner products and tensor powers. For convenience, let us write $[d] = \{1, 2, \ldots, d\}$. Let $x = (x_1, x_2, \ldots, x_d)$ and $y = (y_1, y_2, \ldots, y_d)$ be two vectors of indeterminates. The inner product

$$\langle x, y \rangle = x_1 y_1 + x_2 y_2 + \ldots + x_d y_d \tag{7}$$

is a multilinear polynomial in the indeterminates x, y. For a (multi)set of indices $I \subseteq [d]$ we use the notation

$$\langle x, y\rangle_I = \sum_{i\in I} x_i y_i$$

for an inner product $\langle x, y\rangle$ taken over the indices given by I.

For a positive integer p, we observe the pth power of the inner product $\langle x, y\rangle$ is a polynomial in x, y whose monomials can be indexed by p-tuples $\vec{i} = (i_1, i_2, \ldots, i_p) \in [d]^p$ with $i_1, i_2, \ldots, i_p \in [d]$. Indeed, expanding the product of sums into a sum of products, we observe that

$$\langle x, y\rangle^p = (x_1y_1 + x_2y_2 + \ldots + x_dy_d)^p = \sum_{\vec{i}\in[d]^p} \prod_{\ell=1}^{p} x_{i_\ell} \prod_{\ell=1}^{p} y_{i_\ell}\,. \tag{8}$$

Let us write $x^{\otimes p}$ and $y^{\otimes p}$ for the d^p-dimensional vectors obtained by taking the pth tensor power of x and y and whose coordinates are indexed by the d^p possible p-tuples:

$$x^{\otimes p} = \left(\prod_{\ell=1}^{p} x_{i_\ell} : \vec{i} \in [d]^p\right) \qquad \text{and} \qquad y^{\otimes p} = \left(\prod_{\ell=1}^{p} y_{i_\ell} : \vec{i} \in [d]^p\right).$$

With this notation, we observe from (7) and (8) that

$$\langle x, y\rangle^p = \langle x^{\otimes p}, y^{\otimes p}\rangle\,. \tag{9}$$

3.3. **Cartesian sampling lemma.** The following lemma will be convenient when analysing the concentration of samples obtained in the compression algorithm. For an analogous analysis, cf. Drineas, Kannan, and Mahoney [17].

Lemma 21 (Cartesian sampling)**.** *Let s and m be positive integer squares and let*

$$x_1, x_2, \ldots, x_{m^{1/2}} \in \{-1, 1\} \qquad \text{and} \qquad y_1, y_2, \ldots, y_{m^{1/2}} \in \{-1, 1\}$$

with

$$\sum_{u=1}^{m^{1/2}} x_u = \xi m^{1/2} \qquad \text{and} \qquad \sum_{v=1}^{m^{1/2}} y_v = \eta m^{1/2}$$

for some $-1 \le \xi, \eta \le 1$. Suppose S_1, S_2 are two multisets of size $s^{1/2}$ selected by drawing two $s^{1/2}$-tuples consisting of elements of $[m^{1/2}]$ independently and uniformly at random (with repetition). Then, with high probability in s, we have

$$\left|\sum_{(u,v)\in S_1\times S_2} x_u y_v\right| \le |\xi\eta| s + (|\xi| + |\eta|)s^{3/4}\log s + s^{1/2}(\log s)^2\,. \tag{10}$$

Furthermore, if

$$|\xi|, |\eta| \ge s^{-1/4}\log s\,, \tag{11}$$

then, with high probability in s, we have

$$\left|\sum_{(u,v)\in S_1\times S_2} x_u y_v\right| \ge |\xi\eta| s - (|\xi| + |\eta|)s^{3/4}\log s + s^{1/2}(\log s)^2\,. \tag{12}$$

Proof. For $j = 1, 2, \ldots, s^{1/2}$ let $X_j \in \{-1, 1\}$ be a random variable that independently assumes each of the $m^{1/2}$ values x_u with probability $m^{-1/2}$. For $j = 1, 2, \ldots, s^{1/2}$ let $Y_j \in \{-1, 1\}$ be a random variable that independently assumes each of the $m^{1/2}$ values y_v with probability $m^{-1/2}$. Consider the random variables $X = X_1 + X_2 + \ldots + X_{s^{1/2}}$ and $Y = Y_1 + Y_2 + \ldots + Y_{s^{1/2}}$. By linearity of expectation, we have $\mathrm{E}[X] = \xi s^{1/2}$ and $\mathrm{E}[Y] = \eta s^{1/2}$. Furthermore, from the Hoeffding bound (6) it follows that $|X - \mathrm{E}[X]| \leq s^{1/4} \log s$ and $|Y - \mathrm{E}[Y]| \leq s^{1/4} \log s$ with high probability in s. Thus, we conclude that with high probability in s it holds that

$$\begin{aligned} |XY| &\leq \big(|\xi| s^{1/2} + s^{1/4} \log s\big)\big(|\eta| s^{1/2} + s^{1/4} \log s\big) \\ &= |\xi\eta| s + s^{3/4}(|\xi| + |\eta|) \log s + s^{1/2} (\log s)^2 \,. \end{aligned}$$

Similarly, assuming that $|\xi|, |\eta| \geq s^{-1/4} \log s$, with high probability in s we have the lower bound

$$\begin{aligned} |XY| &\geq \big(|\xi| s^{1/2} - s^{1/4} \log s\big)\big(|\eta| s^{1/2} - s^{1/4} \log s\big) \\ &= |\xi\eta| s - s^{3/4}(|\xi| + |\eta|) \log s + s^{1/2} (\log s)^2 \,. \end{aligned}$$

□

3.4. **Anti-concentration of signed aggregation.** We recall the following anti-concentration lemma from the analysis Valiant's algorithm.

Lemma 22 (Valiant [53, Lemma 3.2])**.** *Let C be a $t \times t$ matrix with entries c_{ij}. Suppose that $\alpha_1, \alpha_2, \ldots, \alpha_t \in \{-1, 1\}$ and $\beta_1, \beta_2, \ldots, \beta_t \in \{-1, 1\}$ have been selected independently and uniformly at random. Then,*

$$\Pr\left(\left|\sum_{i,j} \alpha_i \beta_j c_{ij}\right| \geq \frac{1}{4} \max_{i,j} |c_{ij}|\right) \geq \frac{1}{4} \,.$$

4. The algorithm

This section proves Theorem 6. We present a self-contained analysis and description of the entire algorithm underlying Theorem 6 even if our main technical contribution occurs in the compression subroutine that simultaneously expands and aggregates (cf. §1.5 and §1.4).

4.1. **The input and the parameters.** Let $0 < \gamma < 1$ and $\Delta \geq 1$ be fixed constants. Suppose we are given as input two $d \times n$ matrices $A, B \in \{-1, 1\}^{d \times n}$ together with parameters ρ, τ that satisfy $n^{-\Theta(1)} \leq \tau < \rho < 1$ and $\log_\tau \rho \leq 1 - \Delta^{-1}$. For $i \in [d]$ and $j \in [n]$, let us write a_{ij} and b_{ij} for the entries of A and B at row i, column j. Furthermore, let us write a_j and b_j for the jth column of A and B, respectively.

The algorithm works with three positive integer parameters s, t, p whose precise values we will fix in what follows. At this point we can assume that p is even, that t divides n, and that s is a positive integer square.[13]

[13]To ensure that our eventual choice of t divides n, without loss of generality we can pad the matrices A and B with at most $t - 1$ columns full of 0s so that t divides n. Indeed, each column of 0s will have zero inner product with all the other columns and hence will not affect the behavior of the algorithm.

4.2. **The algorithm.** We describe the algorithm first and then proceed to analyse its correctness and running time. The algorithm executes a five-phase iteration $\lceil (\log n)^2 \rceil$ times, and gives as its output the union of the outputs of all the iterations. The five phases are as follows:

0. Setup. Draw uniformly at random a partition $J_1, J_2, \ldots, J_{n/t}$ of $[n]$ into sets J_k with $|J_k| = t$ for all $k \in [n/t]$. Draw independently and uniformly at random (with repetition) two $s^{1/2}$-tuples consisting of elements of $[d]^{p/2}$ and let I_1 and I_2 be the resulting multisets of size $s^{1/2}$. Form the Cartesian product $I = I_1 \times I_2$ of size s and observe that all the elements of I are p-tuples in $[d]^p$. Draw $\alpha_1, \alpha_2, \ldots, \alpha_n \in \{-1, 1\}$ and $\beta_1, \beta_2, \ldots, \beta_n \in \{-1, 1\}$ independently and uniformly at random.

The compressed matrices (definition). Let us first define what we compute and only then give the algorithm that computes the matrices. For $\vec{i} \in I$ and $k \in [n/t]$, define

$$\hat{a}_{\vec{i}k} = \sum_{j \in J_k} \alpha_j (a_j^{\otimes p})_{\vec{i}} \qquad \text{and} \qquad \hat{b}_{\vec{i}k} = \sum_{j \in J_k} \beta_j (b_j^{\otimes p})_{\vec{i}}. \tag{13}$$

This defines the matrices $\hat{A}_I = (\hat{a}_{\vec{i}k})$ and $\hat{B}_I = (\hat{a}_{\vec{i}k})$, both of size $s \times (n/t)$.[14]

1+2. Compression by simultaneous expansion and aggregation. We compute the compressed matrices $\hat{A}_I$ and $\hat{B}_I$ from the given input A and B as follows. We describe the computation only for $\hat{A}_I$, the computation for $\hat{B}_I$ is symmetric. Let us first establish that the task of computing $\hat{A}_I$ amounts to n/t matrix products *because our random sample I decomposes into the Cartesian product $I = I_1 \times I_2$.* Indeed, we observe from (13) that if we write $\vec{i} = (\vec{i}_1, \vec{i}_2) \in I$ in terms of its parts $\vec{i}_1 \in I_1$ and $\vec{i}_2 \in I_2$, we have

$$\hat{a}_{\vec{i}k} = \hat{a}_{(\vec{i}_1, \vec{i}_2), k} = \sum_{j \in J_k} \alpha_j (a_j^{\otimes p/2})_{\vec{i}_1} (a_j^{\otimes p/2})_{\vec{i}_2}\,. \tag{14}$$

Thus, for each fixed $k \in [n/t]$ we observe that (14) is in fact a matrix product. Indeed, define the $s^{1/2} \times t$ matrix $L_k = (\ell_{\vec{i}_1 j})$ to consist of the entries

$$\ell_{\vec{i}_1 j} = \alpha_j (a_j^{\otimes p/2})_{\vec{i}_1} \tag{15}$$

for $\vec{i}_1 \in I_1$ and $j \in J_k$. Similarly, define the $s^{1/2} \times t$ matrix $R_k = (r_{\vec{i}_2 j})$ to consist of the entries

$$r_{\vec{i}_2 j} = (a_j^{\otimes p/2})_{\vec{i}_2} \tag{16}$$

for $\vec{i}_2 \in I_2$ and $j \in J_k$. By (13) the $(\vec{i}_1, \vec{i}_2)$-entry of the product matrix $L_k R_k^\top$ is precisely the value $\hat{a}_{(\vec{i}_1, \vec{i}_2), k}$. Indeed, $(L_k R_k^\top)_{\vec{i}_1, \vec{i}_2} = \sum_{j \in J_k} \ell_{\vec{i}_1 j} r_{\vec{i}_2 j} = \sum_{j \in J_k} \alpha_j (a_j^{\otimes p/2})_{\vec{i}_1} (a_j^{\otimes p/2})_{\vec{i}_2} = \hat{a}_{(\vec{i}_1, i_2), k}$. To compute the matrix $\hat{A}_I$, we thus proceed as follows. For each $k \in [n/t]$, we construct the matrices L_k and R_k entrywise using (15) and (16), respectively, and then we compute the $s^{1/2} \times s^{1/2}$ product $L_k R_k^\top$. This gives us the $(n/t) \times s$ matrix $\hat{A}_I$, at the cost

[14] We observe in particular these matrices are essentially what results if we execute phases and 1 and 2 of Valiant's algorithm (recall §1.4), with the difference that our sample I always has Cartesian product structure, whereas Valiant's algorithm uses a uniform random sample of $[d]^p$.

of n/t matrix multiplications with inner dimension t and outer dimension $s^{1/2}$.

3. Detection. We compute the matrix product $\hat{A}_I^\top \hat{B}_I$ and *mark* all entries $(k_1, k_2) \in [n/t] \times [n/t]$ that satisfy $|(\hat{A}_I^\top \hat{B}_I)_{k_1 k_2}| \geq \frac{1}{8}\rho^p s$. (This requires one matrix product with inner dimension s and outer dimension n/t.)

4. Listing. Finally, for each marked entry $(k_1, k_2) \in [n/t] \times [n/t]$, we iterate over all $(j_1, j_2) \in J_{k_1} \times J_{k_2}$, compute the inner product $\langle a_{j_1}, b_{j_2} \rangle$, and output (j_1, j_2) as an outlier pair if $|\langle a_{j_1}, b_{j_2} \rangle| \geq \rho d$.

4.3. **Analysis (correctness).** Let us start by fixing the value of the parameter t. With foresight, let t be the unique integer that satisfies

$$n^\gamma \leq t < n^\gamma + 1\,. \tag{17}$$

We begin our analysis by studying the entries of the product matrix $\hat{A}_I^\top \hat{B}_I$. Let $k_1, k_2 \in [n/t]$ and consider the (k_1, k_2)-entry of $\hat{A}_I^\top \hat{B}_I$. From (13) we have

$$\begin{aligned} \big(\hat{A}_I^\top \hat{B}_I\big)_{k_1 k_2} &= \sum_{\vec{i} \in I} \Big(\sum_{j_1 \in J_{k_1}} \alpha_{j_1} \big(a_{j_1}^{\otimes p}\big)_{\vec{i}} \Big) \Big(\sum_{j_2 \in J_{k_2}} \beta_{j_2} \big(b_{j_2}^{\otimes p}\big)_{\vec{i}} \Big) \\ &= \sum_{j_1 \in J_{k_1}} \alpha_{j_1} \sum_{j_2 \in J_{k_2}} \beta_{j_2} \langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I\,. \end{aligned} \tag{18}$$

The value of an entry of $\hat{A}_I^\top \hat{B}_I$ that aggregates only small partial inner products. Let us first derive an upper bound for $|\big(\hat{A}_I^\top \hat{B}_I\big)_{k_1 k_2}|$ subject to an upper bound for the absolute values of the partial inner products $|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I|$. That is, let us assume $U \geq 0$ is an upper bound $|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I| \leq U$ that holds for all $j_1 \in J_{k_1}$ and $j_2 \in J_{k_2}$ with high probability in n. Recall that $|J_{k_1}| = |J_{k_2}| = t$ and that the signs $\alpha_{j_1}, \beta_{j_2} \in \{-1, 1\}$ have been selected independently and uniformly at random for $j_1 \in J_{k_1}$ and $j_2 \in J_{k_2}$. Let us fix $j_1 \in J_{k_1}$ arbitrarily and analyse the concentration of the innermost sum $Z_{j_1} = \sum_{j_2 \in J_{k_2}} \beta_{j_2} \langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I$ in (18). Because β_{j_2} and $\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I$ are independent, each summand $Z_{j_1, j_2} = \beta_{j_2} \langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I$ is a random variable with zero expectation. Let us condition on our assumption $|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I| \leq U$ and thus conclude that $|Z_{j_1, j_2}| \leq U$ holds for each summand. Thus, since $t \leq n$, the Hoeffding bound (6) gives us that $|Z_{j_1}| \leq t^{1/2} U \log n$ with high probability in n. By the union bound we thus have that $|Z_{j_1}| \leq t^{1/2} U \log n$ holds *for all* $j_1 \in J_{k_1}$ with high probability in n. Let us condition on this event. Observe that (18) equals $Z = \sum_{j_1 \in J_{k_1}} \alpha_{j_1} Z_{j_1}$. Since α_{j_1} and Z_{j_1} are independent, the product $\alpha_{j_1} Z_{j_1}$ has zero expectation. Thus, from $|\alpha_{j_1} Z_{j_1}| \leq t^{1/2} U \log n$ and the Hoeffding bound (6) we conclude that, with high probability in n, we have

$$|\big(\hat{A}_I^\top \hat{B}_I\big)_{k_1 k_2}| \leq t U (\log n)^2\,. \tag{19}$$

The value of an entry of $\hat{A}_I^\top \hat{B}_I$ that aggregates at least one large partial inner product. Next, let us suppose that $L \geq 0$ is a lower bound $|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p} \rangle_I| \geq L$ for some specific $j_1, j_2 \in [n]$. Then, for the pair of blocks

$(k_1, k_2) \in [n/t] \times [n/t]$ with $j_1 \in J_{k_1}$ and $j_2 \in J_{k_2}$ we conclude from (18) and Lemma 22 that, with probability at least $1/4$, we have

$$|\big(\hat{A}_I^\top \hat{B}_I\big)_{k_1 k_2}| \geq \frac{L}{4}\,. \tag{20}$$

When (20) holds we say the pair (j_1, j_2) is *signalled by* (the (k_1, k_2)-entry in) $\hat{A}_I^\top \hat{B}_I$.

Concentration of the partial inner products. Let $j_1, j_2 \in [n]$. Let us now study the concentration of $\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_I$ using the Cartesian concentration lemma (Lemma 21). Because $I = I_1 \times I_2$, the partial inner product $\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_I$ decomposes into

$$\begin{aligned}\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_{I_1\times I_2} &= \sum_{\vec{i}\in I_1\times I_2} \big(a_{j_1}^{\otimes p}\big)_{\vec{i}}\big(b_{j_2}^{\otimes p}\big)_{\vec{i}} \\ &= \sum_{(\vec{i}_1,\vec{i}_2)\in I_1\times I_2} \big(a_{j_1}^{\otimes p/2}\big)_{\vec{i}_1}\big(a_{j_1}^{\otimes p/2}\big)_{\vec{i}_2}\big(b_{j_2}^{\otimes p/2}\big)_{\vec{i}_1}\big(b_{j_2}^{\otimes p/2}\big)_{\vec{i}_2}\,.\end{aligned} \tag{21}$$

Now suppose that $\langle a_{j_1}, b_{j_2}\rangle = d\sigma$ for some $-1 \leq \sigma \leq 1$. From (9) we thus have $\langle a_{j_1}^{\otimes p/2}, b_{j_2}^{\otimes p/2}\rangle = d^{p/2}\sigma^{p/2}$. Recalling that $|I_1| = |I_2| = s^{1/2}$, take $\xi = \eta = \sigma^{p/2}$, $m = d^p$, $S_1 = I_1$, $S_2 = I_2$, $x_{\vec{u}} = (a_{j_1}^{\otimes p/2})_{\vec{u}}(b_{j_2}^{\otimes p/2})_{\vec{u}}$, and $y_{\vec{v}} = (a_{j_1}^{\otimes p/2})_{\vec{v}}(b_{j_2}^{\otimes p/2})_{\vec{v}}$ for all $\vec{u}, \vec{v} \in [d]^{p/2}$ in Lemma 21 and thus observe from (21) and (10) that, with high probability in s, we have the upper bound

$$\begin{aligned}\big|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_{I_1\times I_2}\big| &= \Big|\sum_{(\vec{i}_1,\vec{i}_2)\in I_1\times I_2} x_{\vec{i}_1} y_{\vec{i}_2}\Big| \\ &\leq |\sigma^p| s + 2|\sigma^{p/2}| s^{3/4}\log s + s^{1/2}(\log s)^2\,.\end{aligned} \tag{22}$$

Conversely, assuming that

$$|\sigma^p| \geq s^{-1/2}(\log s)^2 \tag{23}$$

so that (11) holds, from (21) and (12) we have, with high probability in s, the lower bound

$$\begin{aligned}\big|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_{I_1\times I_2}\big| &\geq |\sigma^p| s - 2|\sigma^{p/2}| s^{3/4}\log s + s^{1/2}(\log s)^2 \\ &\geq |\sigma^p| s - 2|\sigma^{p/2}| s^{3/4}\log s\,.\end{aligned} \tag{24}$$

Our eventual choice of s will grow at least as fast as a root function of n, so by the union bound we can assume that the upper bound (22) and the lower bound (24) hold for all relevant $j_1, j_2 \in [n]$ with high probability in n.

An upper bound for aggregated background correlations. Let us now fix s to be the least integer square that satisfies[15]

$$\tau^{-2p} \leq s \leq 2\tau^{-2p}\,. \tag{25}$$

Thus, from (22) with $|\sigma| \leq \tau$ we have

$$\big|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_{I_1\times I_2}\big| \leq 4\tau^p s(\log s)^2 \leq 4\tau^p s(\log n)^3\,; \tag{26}$$

[15]Observe that for all large enough positive real x the interval $[x, 2x]$ contains at least one integer square. Also observe that we have not yet fixed our value of p, but will do so later.

here we assume that our eventual choice for s grows no faster than a polynomial function of n, that is $\log s = O(\log n)$ and hence $\log s \leq (\log n)^{3/2}$ for all large enough n. Let us now choose a value for the upper bound U in (19). By (26) we can select

$$U = 8\tau^p s(\log n)^3$$

to conclude that, with high probability n, for all $k_1, k_2 \in [n/t]$ we have

$$|\big(\hat{A}_I^\top \hat{B}_I\big)_{k_1 k_2}| \leq 8\tau^p st(\log n)^5 \tag{27}$$

unless there is at least one pair $(j_1, j_2) \in J_{k_1} \times J_{k_2}$ with $|\langle a_{j_1}, b_{j_2}\rangle| > \tau d$.

Any fixed outlier correlation will be listed with probability at least $1/4$ *during any fixed iteration of the algorithm.* For any outlier correlation we have $|\sigma| \geq \rho$. Let us assume that our eventual choice of p will be such that for all large enough p we have both

$$\rho^p \geq \tau^p (\log s)^2 \geq s^{-1/2}(\log s)^2 \tag{28}$$

and

$$2|\sigma^{p/2}|s^{3/4}\log s \leq 2|\sigma^{p/2}|\tau^{p/2} s \log s \leq \frac{1}{2}|\sigma^p| s\,. \tag{29}$$

Subject to these assumptions, we observe that (23) holds by (28) and from (24) and (29) we have

$$\big|\langle a_{j_1}^{\otimes p}, b_{j_2}^{\otimes p}\rangle_{I_1 \times I_2}\big| \geq \frac{1}{2}\rho^p s\,. \tag{30}$$

Let us now choose a value for the lower bound L in (20). By (30) we can select

$$L = \frac{1}{2}\rho^p s$$

to conclude that any fixed $(j_1, j_2) \in [n] \times [n]$ with $|\langle a_{j_1}, b_{j_2}\rangle| \geq \rho d$ is signalled by $\hat{A}_I^\top \hat{B}_I$ with probability at least $1/4$. Recall now that our threshold in the algorithm for marking an entry $(k_1, k_2) \in [n/t] \times [n/t]$ is $|(\hat{A}_I^\top \hat{B}_I)_{k_1 k_2}| \geq \frac{1}{8}\rho^p s = L/4$. Thus, assuming $(j_1, j_2) \in [n] \times [n]$ is signalled by $\hat{A}_I^\top \hat{B}_I$, the pair of blocks $(k_1, k_2) \in [n/t] \times [n/t]$ with $(j_1, j_2) \in J_{k_1} \times J_{k_2}$ will be marked by the algorithm in the detection phase, and hence (j_1, j_2) will be listed as an outlier in the listing phase.

All outlier correlations will be listed with high probability in n. Since the algorithm has at least $(\log n)^2$ independent iterations, and each iteration outputs any fixed outlier correlation with probability at least $1/4$, it follows that this fixed outlier correlation is output with high probability in n. Taking the union bound over the at most $q \leq n^2$ outlier correlations, every outlier correlation is listed with high probability in n.

At most q *pairs of blocks will be marked with high probability in* n. It remains to complete the parameterization of the algorithm so that at most q pairs of blocks $(k_1, k_2) \in [n/t] \times [n/t]$ will be marked with high probability in n. Recall that there are at most q pairs $(j_1, j_2) \in [n] \times [n]$ with $|\langle a_{j_1}, b_{j_2}\rangle| > \tau d$. Thus, recalling (27) and our marking threshold $\frac{1}{8}\rho^p s$, it suffices to set up the parameters so that

$$8\tau^p st(\log n)^5 < \frac{1}{8}\rho^p s\,.$$

For example, we can select the unique even integer p with

$$\frac{\log t + 5\log\log n + \log 128}{\log(\rho/\tau)} \le p < \frac{\log t + 5\log\log n + \log 128}{\log(\rho/\tau)} + 2\,. \tag{31}$$

From (25) and (31) we thus have

$$t^{\frac{2\log(1/\tau)}{\log(\rho/\tau)}} (128\log n)^{\frac{10\log(1/\tau)}{\log(\rho/\tau)}} \le s < 2t^{\frac{2\log(1/\tau)}{\log(\rho/\tau)}} (128\log n)^{\frac{10\log(1/\tau)}{\log(\rho/\tau)}} \tau^{-4}\,. \tag{32}$$

Now recall from (17) that $n^\gamma \le t < 2n^\gamma$ and observe that

$$1 \le \frac{\log(1/\tau)}{\log(\rho/\tau)} = \frac{1}{1 - \log_\tau \rho} \le \Delta\,.$$

Thus, from (32) we get

$$n^{2\gamma}(128\log n)^{10} \le s < 2^{1+2\Delta} n^{2\Delta\gamma} (128\log n)^{10\Delta} \tau^{-4}\,. \tag{33}$$

Since $\Delta \ge 1$ and $0 < \gamma < 1$ are constants, and recalling our assumption $\tau \ge n^{-\Theta(1)}$, both s and t grow at least as fast as a root function of n and at most as fast as a polynomial function of n. Thus, high probabilities in s and t are high probabilities in n, and vice versa. It remains to justify our assumptions (28) and (29). To establish (28), observe that (31) and $t \ge n^\gamma$ imply $(\rho/\tau)^p \ge n^\gamma$. Furthermore, (33) and $\tau \ge n^{-\Theta(1)}$ imply $(\log s)^2 = O((\log n)^2)$. Thus, (28) holds by (25). To establish (29), recall that $|\sigma| \ge \rho$ and hence it suffices to establish $(\rho/\tau)^{p/2} \ge 4\log s$. The reasoning is similar as we used to establish (28). This completes the correctness proof.

4.4. **Analysis (running time).** From (33) and (17) we have $s = \tilde{O}(n^{2\Delta\gamma}\tau^{-4})$ and $t = \Theta(n^\gamma)$. The matrices $\hat{A}_I$ and $\hat{B}_I$ are computed with n/t matrix multiplications, each with outer dimension $s^{1/2} = \tilde{O}(n^{\Delta\gamma}\tau^{-2})$ and inner dimension $t = \Theta(n^\gamma)$. The exponent of running time for computing the matrices $\hat{A}_I$ and $\hat{B}_I$ from the input A and B is thus

$$\begin{aligned} &1 - \gamma + M(\Delta\gamma - 2\log_n \tau, \gamma) \\ &\quad \le 1 - \gamma + M(\Delta\gamma, \gamma) - 4\log_n \tau\,. \end{aligned}$$

The matrix product $\hat{A}_I^\top B_I$ has outer dimension n/t and inner dimension s. Thus, the exponent of running time for multiplying $\hat{A}_I^\top$ and $\hat{B}_I$ is

$$\begin{aligned} &M(1-\gamma, 2\Delta\gamma - 4\log_n \tau) \\ &\quad \le M(1-\gamma, 2\Delta\gamma) - 4\log_n \tau\,. \end{aligned}$$

The running time (2) for approximate detection in Theorem 6 thus follows. (Indeed, the $\tilde{O}(\cdot)$-notation subsumes the polylogarithmic factors resulting from iterating the algorithm.)

To obtain the the running time for listing, by the analysis in §4.3 at most q entries of $\hat{A}_I^\top B_I$ are marked with high probability in n. Each marked entry induces a computation of $t^2 = \Theta(n^{2\gamma})$ inner products of dimension d, which results in the claimed running time $\tilde{O}(qdn^{2\gamma})$.

This completes the proof of Theorem 6. □

5. Corollaries

5.1. **Proof of Corollary 7.** Let us recall the following basic property of matrix multiplication exponents. We have

$$M(\mu,\nu) \leq \begin{cases} (\omega-1)\mu+\nu & \text{if } \mu \leq \nu; \\ 2\mu+(\omega-2)\nu & \text{if } \mu \geq \nu. \end{cases} \tag{34}$$

Let us now parameterize Theorem 6 to obtain Corollary 7. Let us take

$$\gamma = \frac{1}{2\Delta+1}\,.$$

In this case we have

$$1-\gamma = 2\Delta\gamma\,.$$

Recall the two terms in the maximum in (2). Since $\Delta \geq 1$, from (34) we have that the first term of the maximum in (2) is bounded by

$$1-\gamma+M(\Delta\gamma,\gamma) \leq \frac{4\Delta-2+\omega}{2\Delta+1}\,. \tag{35}$$

The second term of the maximum in (2) is

$$M(1-\gamma,2\Delta\gamma) = \frac{2\Delta\omega}{2\Delta+1}\,. \tag{36}$$

In particular, for $2 \leq \omega \leq 3$ and $\Delta \geq 1$ we observe that $2\Delta\omega \geq 4\Delta-2+\omega$, thus (36) dominates (35). Let us take

$$\Delta = \frac{1}{1-\log_\tau \rho}$$

and observe that (36) simplifies to

$$\frac{2\omega}{3-\log_\tau \rho}\,.$$

This establishes the running time for approximate detection in Corollary 7. The running time for listing is immediate by our choice of γ. □

5.2. **Proof of Corollary 8.** Let us take

$$\gamma = \frac{\alpha}{2\Delta+\alpha}\,,$$

where $0.30298 < \alpha \leq 1$ is the exponent for rectangular matrix multiplication [31], and recall that $\omega \leq 3-\alpha$. In particular, from (34) and $\Delta \geq 1$ we thus have

$$\begin{aligned} &1-\gamma+M(\Delta\gamma,\gamma) \\ &\quad\leq 1-\frac{\alpha}{2\Delta+\alpha}+\frac{2\alpha\Delta}{2\Delta+\alpha}+\frac{(\omega-2)\alpha}{2\Delta+\alpha} \\ &\quad= \frac{2\Delta+\alpha-\alpha+\alpha(2\Delta+\omega-2)}{2\Delta+\alpha} \\ &\quad= \frac{2\Delta(\alpha+1)+\alpha(\omega-2)}{2\Delta+\alpha} \\ &\quad\leq \frac{2\Delta(\alpha+1)+\alpha(1-\alpha)}{2\Delta+\alpha} \\ &\quad\leq \frac{4\Delta}{2\Delta+\alpha}\,. \end{aligned} \tag{37}$$

Furthermore, by definition of α and our choice of γ, we have

$$M(1-\gamma, 2\Delta\gamma) = M\left(\frac{2\Delta}{2\Delta+\alpha}, \frac{2\Delta\alpha}{2\Delta+\alpha}\right) = \frac{4\Delta}{2\Delta+\alpha}\,. \tag{38}$$

Thus (38) dominates (37). Simplifying (38), we obtain

$$\frac{4}{2+\alpha\big(1-\log_\tau \rho\big)}\,.$$

This establishes the running time for approximate detection in Corollary 8. The running time for listing is immediate by our choice of γ. $\square$

5.3. **Proof of Corollary 12.** Fix a constant $0 < \epsilon < \omega/3$ and let $n^{-\Theta(1)} < \rho < 1$ be given. With the objective of eventually applying Theorem 6, let us begin by selecting a suitable value of τ. First, we want a small enough $0 < \delta < 1$ so that

$$\frac{2\omega}{3-\delta} \le \frac{2\omega}{3} + \epsilon\,,$$

or what is the same,

$$\delta \le \frac{9\epsilon}{2\omega+3\epsilon} = \frac{9}{2\omega/\epsilon+3}\,. \tag{39}$$

To obtain $\log_\tau \rho \le \delta$, let us take $\tau = \rho^{1/\delta}$. Because $0 < \delta < 1$ and $0 < \rho < 1$, we have $\tau < \rho$. Let us now check that for all large enough d all the inner products between pairs of vectors in the input, except for the planted pair, are at most τd in absolute value with probability $1 - o(1)$. For

$$d \ge 5\rho^{-2/\delta}\log n$$

it follows from the Hoeffding bound (6) and the union bound that with probability $1 - o(1)$ all of the at most n^2 pairs of vectors that have at least one independent uniform random vector in the pair have inner product at most τd in absolute value. Indeed, as $n \to \infty$ we have

$$2n^2 \exp\left(-\frac{2(\tau d)^2}{4d}\right) = 2n^2 \exp\left(-\frac{\tau^2 d}{2}\right) \le 2n^2 \exp\left(-\frac{5}{2}\log n\right) \to 0\,.$$

It remains to find the planted pair with inner product at least ρd in absolute value. Let us apply the approximate detection algorithm in Theorem 6. To parameterize the algorithm, take $\Delta = 1/(1-\delta)$ and $\gamma = 1/(2\Delta+1)$. Observe in particular that both Δ and γ depend only on the constants δ and ϵ, and do not depend on ρ or τ. Mimic the analysis in the proof of Corollary 7 to obtain from (2) the running time bound

$$\tilde{O}(n^{2\omega/3+\epsilon}\tau^{-4})\,. \tag{40}$$

Observe furthermore that within the same time bound (40) we can run the approximate detection algorithm recursively $O(\log\log n)$ times on the pair of blocks marked by the algorithm to find the planted pair, w.h.p. We obtain (3) by combining (40) with $\tau = \rho^{1/\delta}$ and (39). $\square$

5.4. **Proof of Corollary 14.** This proof relies on a split-and-list idea outlined by Valiant [53, p. 32]; we present a proof here for completeness of exposition.

Let us start by setting up some notation. For a subset $A \subseteq [n]$ and a vector $(x_1, x_2, \ldots, x_n) \in \{-1,1\}^n$, let us write $x_A = \prod_{\ell \in A} x_\ell$. (For the empty set, define $x_\emptyset = 1$.) Observe that for all $A, B \subseteq [n]$ it holds that $x_A x_B = x_{A \oplus B}$, where $A \oplus B = (A \setminus B) \cup (B \setminus A)$ is the symmetric difference of A and B. Let us write $\binom{[n]}{k}$ for the set of all k-subsets of $[n]$.

Fix a constant $0 < \epsilon < \omega/3$. Select a corresponding constant $0 < \delta < 1$ so that (39) holds.

Let $0 < \eta < 1$ be the given noise rate with $n^{-\Theta(1)} < |1 - 2\eta| < 1$. Let $S \subseteq [n]$ be the support of the parity function, $|S| = k \geq 2$. Our task is to determine S by drawing examples. Let us draw d examples $(x, y) \in \{-1,1\}^d \times \{-1,1\}$. (We will fix a lower bound for d in what follows.) Recall that each label y in an example has the structure $y = z x_S$, where $z \in \{-1,1\}$ independently with $\Pr(z = -1) = \eta$.

Construct a collection of $\binom{n}{\lfloor k/2 \rfloor} + \binom{n}{\lceil k/2 \rceil} \leq 2n^{(k+1)/2}$ vectors of dimension d as follows. First, for each $J_1 \in \binom{n}{\lfloor k/2 \rfloor}$, construct the vector a_{J_1} whose entries are the values x_{J_1}, where x ranges over the d examples (x, y) that we have drawn. Second, for each $J_2 \in \binom{n}{\lceil k/2 \rceil}$, construct the vector b_{J_2} whose entries are the values $x_{J_2} y$, where x and y range over the d examples (x, y) that we have drawn.

Let us write $\mathrm{Bin}_{\pm 1}(d, \beta)$ for the sum of d random variables, each independently taking values in $\{-1, 1\}$ such that β is the probability of taking the value -1.

Let us study the inner products between the vectors in our collection. We will use the notation "$\sum_{(x,y)}$" to indicate a sum over the d examples (x, y) that we have drawn. Observe that for all $J_1, J_1' \in \binom{n}{\lfloor k/2 \rfloor}$ we have

$$\langle a_{J_1}, a_{J_1'} \rangle = \sum_{(x,y)} x_{J_1} x_{J_1'} = \sum_{(x,y)} x_{J_1 \oplus J_1'} \,. \tag{41}$$

Similarly, for all $J_2, J_2' \in \binom{n}{\lceil k/2 \rceil}$ we have

$$\langle b_{J_2}, b_{J_2'} \rangle = \sum_{(x,y)} x_{J_2} y x_{J_2'} y = \sum_{(x,y)} x_{J_2 \oplus J_2'} \,. \tag{42}$$

Finally, for all $J_1 \in \binom{n}{\lfloor k/2 \rfloor}$ and $J_2 \in \binom{n}{\lceil k/2 \rceil}$ we have

$$\langle a_{J_1}, b_{J_2} \rangle = \sum_{(x,y)} x_{J_1} x_{J_2} y = \sum_{(x,y)} x_{J_1} x_{J_2} x_S z = \sum_{(x,y)} x_{J_1 \oplus J_2 \oplus S} z \,. \tag{43}$$

Recalling that the vector x in each example is an independent uniform random vector in $\{-1,1\}^n$, we observe[16] from (41), (42), and (43) that each of the following inner products has distribution $\mathrm{Bin}_{\pm 1}(d, 1/2)$: (i) $\langle a_{J_1}, a_{J_1'} \rangle$ for distinct J_1, J_1', (ii) $\langle b_{J_2}, b_{J_2'} \rangle$ for distinct J_2, J_2', and (iii) $\langle a_{J_1}, b_{J_2} \rangle$ for all

[16] Indeed, for all *nonempty* $K \subseteq [n]$, the random variables $\sum_{(x,y)} x_K$ and $\sum_{(x,y)} x_K z$ have distribution $\mathrm{Bin}_{\pm 1}(d, 1/2)$. For the symmetric difference it holds that $K_1 \oplus K_2 = \emptyset$ if and only if $K_1 = K_2$.

J_1, J_2 with $J_1 \oplus J_2 \neq S$. Furthermore, all the *remaining* inner products between distinct vectors in our collection, that is, $\langle a_{J_1}, b_{J_2} \rangle$ for all J_1, J_2 with $J_1 \oplus J_2 = S$, have distribution $\mathrm{Bin}_{\pm 1}(d, \eta)$. This difference between distributions enables us to detect a pair J_1, J_2 with $J_1 \oplus J_2 = S$ and hence the set S.

Take $\rho = |1 - 2\eta|$ and observe that $n^{-\Theta(1)} < \rho < 1$. Let us take $\tau = \rho^{1/\delta}$. Because $0 < \delta < 1$ and $0 < \rho < 1$, we have $\tau < \rho$. Furthermore, $\log_\tau \rho \leq \delta$. For

$$d \geq (2k+3)\rho^{-2/\delta} \log n \tag{44}$$

it follows from the Hoeffding bound (6) and the union bound that with probability $1 - o(1)$ all of the at most $2n^{k+1}$ inner products with distribution $\mathrm{Bin}_{\pm 1}(d, 1/2)$ are at most τd in absolute value. Indeed, as $n \to \infty$ we have

$$\begin{aligned} 4n^{k+1} \exp\left(-\frac{2(\tau d)^2}{4d}\right) &= 4n^{k+1} \exp\left(-\frac{\tau^2 d}{2}\right) \\ &\leq 4n^{k+1} \exp\left(-\frac{2k+3}{2} \log n\right) \to 0\,. \end{aligned}$$

Observe that the expectation of $\mathrm{Bin}_{\pm 1}(d, \eta)$ is $(1 - 2\eta)d$ and that the absolute value of the expectation is ρd. For $\eta \geq 1/2$ (respectively, for $\eta \leq 1/2$) the probability that $\mathrm{Bin}_{\pm 1}(d, \eta)$ is at most (respectively, at least) its expectation is at least $1/4$ [21, Theorem 1 and Corollary 3].

Thus, with probability at least $1/4 - o(1)$ over the d examples drawn, both (a) the absolute value of at least one inner product with distribution $\mathrm{Bin}_{\pm 1}(d, \eta)$ is at least ρd, and (b) all of the at most $2n^{k+1}$ inner products with distribution $\mathrm{Bin}_{\pm 1}(d, 1/2)$ are at most τd in absolute value.

Let us apply the approximate detection algorithm in Theorem 6. To parameterize the algorithm, take $\Delta = 1/(1-\delta)$ and $\gamma = 1/(2\Delta+1)$. Observe in particular that both Δ and γ depend only on the constants δ and ϵ, and do not depend on ρ or τ. Mimic the analysis in the proof of Corollary 7 to obtain from (2) for all large enough k the running time bound

$$\tilde{O}(n^{(\omega+\epsilon)k/3}\tau^{-4})\,. \tag{45}$$

Observe furthermore that within the same time bound (45) we can run the approximate detection algorithm recursively $O(\log \log n)$ times on any one pair of blocks marked by the algorithm to find at least one inner product $\langle a_{J_1}, b_{J_2} \rangle$ with absolute value at least ρd, and hence the set $S = J_1 \oplus J_2$, w.h.p. Here we have of course conditioned on both events (a) and (b) above.

Suppose now that (a) fails. Then (b) still happens with probability $1 - o(1)$, so the approximate detection algorithm will, w.h.p., mark no pair of blocks consisting only of inner products with distribution $\mathrm{Bin}_{\pm 1}(d, 1/2)$. Thus we can detect that the algorithm fails (no pair of blocks is marked) and draw more examples as necessary to boost the success probability of the algorithm.

We obtain (5) by combining (45) with $\tau = \rho^{1/\delta}$ and (39). Similarly, we obtain (4) by combining (44) with (39). □

5.5. **Proof of Theorem 15.** The proof is otherwise identical to the proof of Theorem 6 in §4, with the exception that we modify the listing phase as

follows. For each of the at most n^2 pairs of blocks $k_1, k_2 \in [n/t]$ marked by the detection algorithm, we run the detection algorithm again, with parameter κ replacing the parameter γ, and using the same value for the parameter Δ. Indeed, each signalled pair of blocks can be viewed as an input of size n^γ and dimension d. From these inputs we obtain as output at most q pairs of blocks, each of size $n^{\gamma\kappa}$. Finally, we run the listing algorithm for these at most q blocks. The running time and success probability follow immediately from Theorem 6 and the observation that the number of recursive invocations on inputs of size n^γ is at most $q \leq n^2$, so by the union bound the conclusion holds w.h.p. in n. □

5.6. **Proof of Corollary 16.** Analogous to Corollary 7, take $\gamma = \kappa = 1/(2\Delta + 1)$. □

5.7. **Proof of Corollary 17.** Analogous to Corollary 8, take $\gamma = \kappa = \alpha/(2\Delta + \alpha)$. □

6. A lower bound via orthogonal vectors

This section presents local transformation from the orthogonal vectors problem (cf. [1, 60]) to Bichromatic Outlier Correlations. We present this transformation for completeness of exposition only and remark that it has been superseded by very recent results of Ahle, Pagh, Razenshteyn, and Silvestri [3]. First, let us recall the orthogonal vectors problem:

Problem 23 (Orthogonal Vectors)**.** Given as input a set of n vectors with dimension d and entries in $\{0, 1\}$, decide whether there is an orthogonal pair of vectors over the integers.

The sparsification lemma of Impagliazzo, Paturi, and Zane [25] and a lemma of Williams [59] yield the following conditional hardness result (cf. [60, Lemma A.1]) for Orthogonal Vectors:

Lemma 24 (Williams [59])**.** *Suppose there exists a constant $\delta > 0$ and a randomized algorithm that for all constants $c \geq 1$ solves the* Orthogonal Vectors *problem with $d \leq c \log n$ in time $O(n^{2-\delta})$, w.h.p. Then, the Strong Exponential Time Hypothesis is false.*

The Orthogonal Vectors Conjecture (OVC) states that the algorithm assumed in Lemma 24 does not exist (cf. [1, 60]). Lemma 24 together with the following local transformation shows that Bichromatic Outlier Correlations cannot be solved in subquadratic time for inputs with $\log_\tau \rho$ arbitrarily close to 1 unless both SETH and OVC are false:

Lemma 25. *Suppose that there exists a constant $\delta > 0$ and a randomized algorithm that for all constants $c \geq 1$ solves the* Bichromatic Outlier Correlations *problem for Boolean inputs with $d \leq c \log n$ and $|\rho - \tau| \leq 2/d$ in time $O(n^{2-\delta})$, w.h.p. Then, there exists a constant $\delta' > 0$ and a randomized algorithm that for all constants $c' \geq 1$ solves the* Orthogonal Vectors *problem with $d' \leq c' \log n$ in time $O(n^{2-\delta'})$, w.h.p.*

Proof. Without loss of generality we may work with a bichromatic version of Ortogonal Vectors where our input is a pair of matrices $S, T \in$

$\{0,1\}^{d'\times n}$. We transform the matrices S,T into a pair of matrices $A,B \in \{-1,1\}^{(4d'+1)\times n}$ such that for all $j_1,j_2 \in [n]$ it holds that $|\langle a_{j_1}, b_{j_2}\rangle|$ is above a threshold value if and only if $\langle s_{j_1}, t_{j_2}\rangle = 0$.

The local transformation is given by a function $h : \{0,1\}^2 \to \mathbb{Z}$ that decomposes into two functions $u, v : \{0,1\} \to \{-1,1\}^k$ such that $h(x,y) = \langle u(x), v(y)\rangle$ with $h(0,0) = h(1,0) = h(0,1) = k_1$ for some integer $k_1 > 0$, and $h(1,1) = k_2$ for some integer $k_2 < 0$. One such pair of functions with $k = 4$, $k_1 = 2$, and $k_2 = -2$ is the following:

$$u(x) = \begin{bmatrix} 1 & 1-2x & 1 & 1-2x \end{bmatrix}^\top ,$$
$$v(y) = \begin{bmatrix} 1 & 1 & 1-2y & 2y-1 \end{bmatrix}^\top .$$

Apply the function u (respectively, the function v) to each element of S (respectively, T) to produce a matrix A (respectively, B) of size $4d' \times n$. Finally, append one full row of 1-entries to both matrices A and B.

The matrices A and B yield a $\{-1,1\}$-valued input to Outlier Correlations with the following parameters: the number of vectors is n, the dimension is $d = 4d' + 1$, the outlier parameter is $\rho = (2d'+1)/(4d'+1)$, and the background parameter $\tau = (2d'-1)/(4d'+1)$. In particular, we observe that the inner products $\langle a_{j_1}, b_{j_2}\rangle$ lie in the interval $[-2d'+1, 2d'+1]$, and we have $|\langle a_{j_1}, b_{j_2}\rangle| \geq \rho(4d'+1)$ if and only if $\langle s_{j_1}, t_{j_2}\rangle = 0$. Furthermore, without loss of generality we may assume that the instance A, B has at most $q = O(n^{2-\eta})$ outlier pairs for a constant $\eta > 0$. Indeed, we may test uniformly at random $n^{1.1\eta}$ pairs of vectors for outlier pairs. If we find an outlier pair, we are done; otherwise we can conclude that w.h.p. in n the instance has at most q outlier pairs. In the latter case we run the assumed algorithm for Outlier Correlations to find a pair if it exists. Observe that for our choice of ρ, τ we have $|\rho - \tau| = 2/(4d'+1) = 2/d$. □

Remark 26. We can replace the assumption $|\rho - \tau| \leq 2/d$ in Lemma 25 with

$$\log_\tau \rho \geq \frac{\log \frac{2d'+1}{4d'+1}}{\log \frac{2d'-1}{4d'+1}} = 1 - \frac{\log\left(1 + \frac{4}{d-3}\right)}{\log\left(2 + \frac{6}{d-3}\right)} = 1 - \Theta(1/d)\,.$$

Acknowledgments

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement 338077 "Theory and Practice of Advanced Search and Enumeration." We thank the anonymous referees for their helpful comments and Rasmus Pagh and Thomas Dybdahl Ahle for useful discussions. Work done in part while the second author was visiting the Simons Institute for the Theory of Computing.

References

[1] A. Abboud, R. R. Williams, and H. Yu, *More applications of the polynomial method to algorithm design*, in Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015, P. Indyk, ed., SIAM, 2015, pp. 218–230.

[2] P. Achlioptas, B. Schölkopf, and K. M. Borgwardt, *Two-locus association mapping in subquadratic time*, in Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, CA, USA, August 21-24, 2011, C. Apté, J. Ghosh, and P. Smyth, eds., ACM, 2011, pp. 726–734.

[3] T. D. Ahle, R. Pagh, I. Razenshteyn, and F. Silvestri, *On the complexity of inner product similarity join*, CoRR, abs/1510.02824 (2015).

[4] J. Alman and R. Williams, *Probabilistic polynomials and Hamming nearest neighbors*, in 56th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, October 18-20, 2015, 2015. To appear.

[5] A. Andoni and P. Indyk, *Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions*, in 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings, 2006, pp. 459–468.

[6] A. Andoni, P. Indyk, T. Laarhoven, I. Razenshteyn, and L. Schmidt, *Practical optimal LSH for angular distance*, in Proceedings of the 29th Annual Conference on Neural Information Processing Systems, NIPS 2015, Montréal, Canada, December 7-12, 2015, 2015. To appear.

[7] A. Andoni, P. Indyk, H. L. Nguyen, and I. Razenshteyn, *Beyond locality-sensitive hashing*, in Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, 2014, pp. 1018–1028.

[8] A. Andoni and I. Razenshteyn, *Optimal data-dependent hashing for approximate near neighbors*, in Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, R. A. Servedio and R. Rubinfeld, eds., ACM, 2015, pp. 793–801.

[9] R. J. Bayardo, Y. Ma, and R. Srikant, *Scaling up all pairs similarity search*, in Williamson et al. [61], pp. 131–140.

[10] A. Blum, A. Kalai, and H. Wasserman, *Noise-tolerant learning, the parity problem, and the statistical query model*, J. ACM, 50 (2003), pp. 506–519.

[11] M. Charikar, *Similarity estimation techniques from rounding algorithms*, in Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada, J. H. Reif, ed., ACM, 2002, pp. 380–388.

[12] S. Chaudhuri, V. Ganti, and R. Kaushik, *A primitive operator for similarity joins in data cleaning*, in Proceedings of the 22nd International Conference on Data Engineering, ICDE 2006, 3-8 April 2006, Atlanta, GA, USA, L. Liu, A. Reuter, K. Whang, and J. Zhang, eds., IEEE Computer Society, 2006, p. 5.

[13] X. Chen, G. Lebanon, H. Wang, and M. J. Zaki, eds., *21st ACM International Conference on Information and Knowledge Management, CIKM'12, Maui, HI, USA, October 29 - November 02, 2012*, ACM, 2012.

[14] E. Cohen, M. Datar, S. Fujiwara, A. Gionis, P. Indyk, R. Motwani, J. D. Ullman, and C. Yang, *Finding interesting associations without support pruning*, IEEE Trans. Knowl. Data Eng., 13 (2001), pp. 64–78.

[15] R. R. Curtin, A. G. Gray, and P. Ram, *Fast exact max-kernel search*, in Proceedings of the 13th SIAM International Conference on Data Mining, May 2-4, 2013. Austin, Texas, USA., SIAM, 2013, pp. 1–9.

[16] A. Das, M. Datar, A. Garg, and S. Rajaram, *Google news personalization: scalable online collaborative filtering*, in Williamson et al. [61], pp. 271–280.

[17] P. Drineas, R. Kannan, and M. W. Mahoney, *Fast Monte Carlo algorithms for matrices. I. Approximating matrix multiplication*, SIAM J. Comput., 36 (2006), pp. 132–157.

[18] M. Dubiner, *Bucketing coding and information theory for the statistical high-dimensional nearest-neighbor problem*, IEEE Transactions on Information Theory, 56 (2010), pp. 4166–4179.

[19] V. Feldman, P. Gopalan, S. Khot, and A. K. Ponnuswami, *On agnostic learning of parities, monomials, and halfspaces*, SIAM J. Comput., 39 (2009), pp. 606–645.

[20] F. Galton, *Co-relations and their measurement, chiefly from anthropometric data*, Proceedings of the Royal Society of London, 45 (1888), pp. 135–145.

[21] S. Greenberg and M. Mohri, *Tight lower bound on the probability of a binomial exceeding its expectation*, Statistics and Probability Letters, 86 (2014), pp. 91–98.

[22] E. Grigorescu, L. Reyzin, and S. Vempala, *On noise-tolerant learning of sparse parities and related problems*, in Algorithmic Learning Theory - 22nd International Conference, ALT 2011, Espoo, Finland, October 5-7, 2011. Proceedings, J. Kivinen, C. Szepesvári, E. Ukkonen, and T. Zeugmann, eds., vol. 6925 of Lecture Notes in Computer Science, Springer, 2011, pp. 413–424.

[23] W. Hoeffding, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc., 58 (1963), pp. 13–30.

[24] R. Impagliazzo and R. Paturi, *On the complexity of k-SAT*, J. Comput. Syst. Sci., 62 (2001), pp. 367–375.

[25] R. Impagliazzo, R. Paturi, and F. Zane, *Which problems have strongly exponential complexity?*, J. Comput. Syst. Sci., 63 (2001), pp. 512–530.

[26] P. Indyk, *Nearest neighbors in high-dimensional spaces*, in Handbook of Discrete and Computational Geometry, J. E. Goodman and J. O'Rourke, eds., Discrete Mathematics and its Applications (Boca Raton), Chapman & Hall/CRC, Boca Raton, FL, second ed., 2004, pp. 877–892.

[27] P. Indyk and R. Motwani, *Approximate nearest neighbors: Towards removing the curse of dimensionality*, in Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998, 1998, pp. 604–613.

[28] M. Kapralov, *Smooth tradeoffs between insert and query complexity in nearest neighbor search*, in Proceedings of the 34th ACM Symposium on Principles of Database Systems, PODS 2015, Melbourne, Victoria, Australia, May 31 - June 4, 2015, T. Milo and D. Calvanese, eds., ACM, 2015, pp. 329–342.

[29] N. Koenigstein, P. Ram, and Y. Shavitt, *Efficient retrieval of recommendations in a matrix factorization framework*, in Chen et al. [13], pp. 535–544.

[30] E. Kushilevitz, R. Ostrovsky, and Y. Rabani, *Efficient search for approximate nearest neighbor in high dimensional spaces*, in Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998, 1998, pp. 614–623.

[31] F. Le Gall, *Faster algorithms for rectangular matrix multiplication*, in 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, 2012, pp. 514–523.

[32] ———, *Powers of tensors and fast matrix multiplication*, in International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014, 2014, pp. 296–303.

[33] D. Lee, J. Park, J. Shim, and S. Lee, *An efficient similarity join algorithm with cosine similarity predicate*, in Database and Expert Systems Applications, 21th International Conference, DEXA 2010, Bilbao, Spain, August 30 - September 3, 2010, Proceedings, Part II, P. G. Bringas, A. Hameurlain, and G. Quirchmayr, eds., vol. 6262 of Lecture Notes in Computer Science, Springer, 2010, pp. 422–436.

[34] Y. Low and A. X. Zheng, *Fast top-k similarity queries via matrix compression*, in Chen et al. [13], pp. 2070–2074.

[35] V. Lyubashevsky, *The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem*, in Approximation, Randomization and Combinatorial Optimization, Algorithms and Techniques, 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2005 and 9th International Workshop on Randomization and Computation, RANDOM 2005, Berkeley, CA, USA, August 22-24, 2005, Proceedings, C. Chekuri, K. Jansen, J. D. P. Rolim, and L. Trevisan, eds., vol. 3624 of Lecture Notes in Computer Science, Springer, 2005, pp. 378–389.

[36] A. May and I. Ozerov, *On computing nearest neighbors with applications to decoding of binary linear codes*, in Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I, 2015, pp. 203–228.

[37] S. Meiser, *Point location in arrangements of hyperplanes*, Inf. Comput., 106 (1993), pp. 286–303.

[38] M. L. Minsky and S. A. Papert, *Perceptrons: An Introduction to Computational Geometry*, MIT Press, Boston, MA, 1969.

[39] E. Mossel, R. O'Donnell, and R. A. Servedio, *Learning functions of k relevant variables*, J. Comput. Syst. Sci., 69 (2004), pp. 421–434.

[40] R. Motwani, A. Naor, and R. Panigrahy, *Lower bounds on locality sensitive hashing*, SIAM J. Discrete Math., 21 (2007), pp. 930–935.

[41] B. Neyshabur and N. Srebro, *On symmetric and asymmetric LSHs for inner product search*, in Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015, F. R. Bach and D. M. Blei, eds., vol. 37 of JMLR Proceedings, JMLR.org, 2015, pp. 1926–1934.

[42] R. O'Donnell, Y. Wu, and Y. Zhou, *Optimal lower bounds for locality-sensitive hashing (except when q is tiny)*, TOCT, 6 (2014), p. 5.

[43] R. Pagh, *Compressed matrix multiplication*, TOCT, 5 (2013), p. 9.

[44] ———, *Locality-sensitive hashing without false negatives*, in Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, Virginia, USA, January 10-12, 2016, 2016. To appear.

[45] R. Paturi, S. Rajasekaran, and J. H. Reif, *The light bulb problem*, in Proceedings of the Second Annual Workshop on Computational Learning Theory, COLT 1989, Santa Cruz, CA, USA, July 31 - August 2, 1989, 1989, pp. 261–268.

[46] K. Pearson, *Notes on the history of correlation*, Biometrika, 13 (1920), pp. 25–45.

[47] P. Ram and A. G. Gray, *Maximum inner-product search using cone trees*, in The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012, Q. Yang, D. Agarwal, and J. Pei, eds., ACM, 2012, pp. 931–939.

[48] A. Shrivastava and P. Li, *Asymmetric LSH (ALSH) for sublinear time maximum inner product search (MIPS)*, in Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds., 2014, pp. 2321–2329.

[49] ———, *Asymmetric minwise hashing for indexing binary inner products and set containment*, in Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015, A. Gangemi, S. Leonardi, and A. Panconesi, eds., ACM, 2015, pp. 981–991.

[50] Y. N. Silva, W. G. Aref, and M. H. Ali, *The similarity join database operator*, in Proceedings of the 26th International Conference on Data Engineering, ICDE 2010, March 1-6, 2010, Long Beach, California, USA, F. Li, M. M. Moro, S. Ghandeharizadeh, J. R. Haritsa, G. Weikum, M. J. Carey, F. Casati, E. Y. Chang, I. Manolescu, S. Mehrotra, U. Dayal, and V. J. Tsotras, eds., IEEE, 2010, pp. 892–903.

[51] S. M. Stigler, *Francis Galton's account of the invention of correlation*, Statistical Science, 4 (1989), pp. 73–86.

[52] C. Teflioudi, R. Gemulla, and O. Mykytiuk, *LEMP: fast retrieval of large entries in a matrix product*, in Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015, T. Sellis, S. B. Davidson, and Z. G. Ives, eds., ACM, 2015, pp. 107–122.

[53] G. Valiant, *Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem*, J. ACM, 62 (2015), p. 13.

[54] L. G. Valiant, *Functionality in neural nets*, in Proceedings of the First Annual Workshop on Computational Learning Theory, COLT '88, Cambridge, MA, USA, August 3-5, 1988, 1988, pp. 28–39.

[55] J. Wang, G. Li, and J. Feng, *Fast-join: An efficient method for fuzzy token matching based string similarity join*, in Proceedings of the 27th International Conference on Data Engineering, ICDE 2011, April 11-16, 2011, Hannover, Germany, S. Abiteboul, K. Böhm, C. Koch, and K. Tan, eds., IEEE Computer Society, 2011, pp. 458–469.

[56] ———, *Can we beat the prefix filtering?: an adaptive framework for similarity join and search*, in Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2012, Scottsdale, AZ, USA, May 20-24, 2012, K. S. Candan, Y. Chen, R. T. Snodgrass, L. Gravano, and A. Fuxman, eds., ACM, 2012, pp. 85–96.

[57] Y. Wang, A. Metwally, and S. Parthasarathy, *Scalable all-pairs similarity search in metric spaces*, in The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, Chicago, IL, USA, August 11-14, 2013, I. S. Dhillon, Y. Koren, R. Ghani, T. E. Senator, P. Bradley, R. Parekh, J. He, R. L. Grossman, and R. Uthurusamy, eds., ACM, 2013, pp. 829–837.

[58] R. Weber, H. Schek, and S. Blott, *A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces*, in VLDB'98, Proceedings of 24rd International Conference on Very Large Data Bases, August 24-27, 1998, New York City, New York, USA, 1998, pp. 194–205.

[59] R. Williams, *A new algorithm for optimal 2-constraint satisfaction and its implications*, Theor. Comput. Sci., 348 (2005), pp. 357–365.

[60] R. Williams and H. Yu, *Finding orthogonal vectors in discrete structures*, in Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, C. Chekuri, ed., SIAM, 2014, pp. 1867–1877.

[61] C. L. Williamson, M. E. Zurko, P. F. Patel-Schneider, and P. J. Shenoy, eds., *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, ACM, 2007.

[62] C. Xia, H. Lu, B. C. Ooi, and J. Hu, *Gorder: An efficient method for KNN join processing*, in (e)Proceedings of the Thirtieth International Conference on Very Large Data Bases, Toronto, Canada, August 31 - September 3 2004, M. A. Nascimento, M. T. Özsu, D. Kossmann, R. J. Miller, J. A. Blakeley, and K. B. Schiefer, eds., Morgan Kaufmann, 2004, pp. 756–767.

[63] C. Xiao, W. Wang, X. Lin, J. X. Yu, and G. Wang, *Efficient similarity joins for near-duplicate detection*, ACM Trans. Database Syst., 36 (2011), p. 15.

[64] R. B. Zadeh and A. Goel, *Dimension independent similarity computation*, Journal of Machine Learning Research, 14 (2013), pp. 1605–1626.

[65] P. Zezula, G. Amato, V. Dohnal, and M. Batko, *Similarity Search. The Metric Space Approach*, vol. 32 of Advances in Database Systems, Springer, New York, NY, 2006.

Department of Computer Science, Aalto University, Helsinki, Finland
E-mail address: matti.karppa@aalto.fi
E-mail address: petteri.kaski@aalto.fi
E-mail address: jukka.kohonen@aalto.fi